**Exam** : **JN0-351**

**Title** : Enterprise Routing and Switching, Specialist (JNCIS-ENT)

**https://www.passcert.com/JN0-351.html**

1.Exhibit.

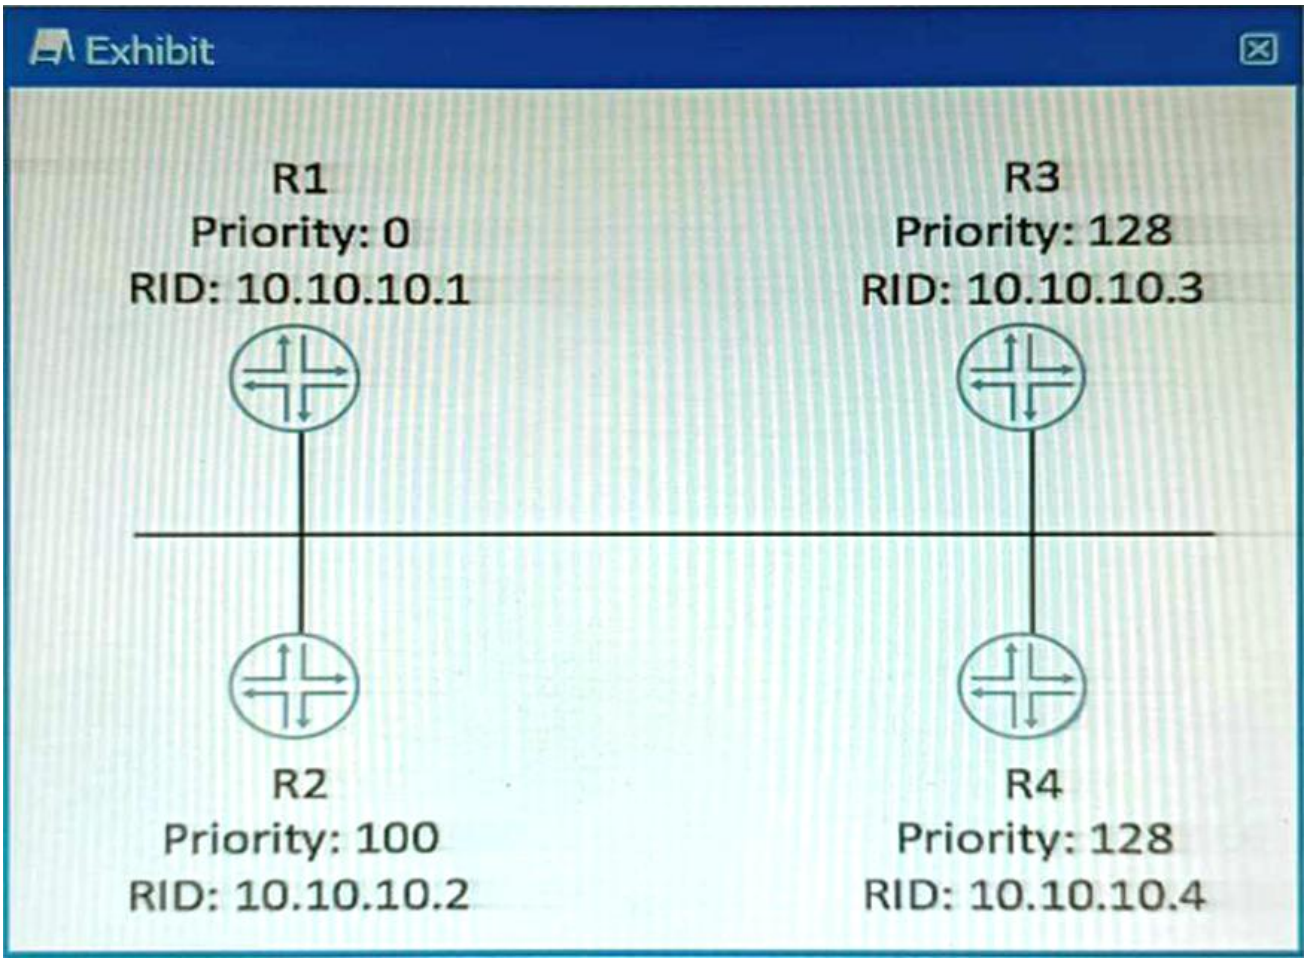


Which router will become the OSPF BDR if all routers are powered on at the same time?

A. R4

B. R1

C. R3

D. R2

**Answer:** A

**Explanation:**

OSPF DR/BDR election is a process that occurs on multi-access data links. It is intended to select two OSPF nodes: one to be acting as the Designated Router (DR), and another to be acting as the Backup Designated Router (BDR). The DR and BDR are responsible for generating network LSAs for the multi-access network and synchronizing the LSDB with other routers on the same network1.

The DR/BDR election is based on two criteria: the OSPF priority and the router ID. The OSPF priority is a value between 0 and 255 that can be configured on each interface participating in OSPF. The default priority is 1. A priority of 0 means that the router will not participate in the election and will never become a DR or BDR. The router with the highest priority will become the DR, and the router with the second highest priority will become the BDR. If there is a tie in priority, then the router ID is used as a tie-breaker. The router ID is a 32-bit number that uniquely identifies each router in an OSPF domain. It can be manually configured or automatically derived from the highest IP address on a loopback interface or any active interface2.

In this scenario, all routers have the same priority of 1, so the router ID will determine the outcome of the

election. The router IDs are shown in the exhibit as RID values. The highest RID belongs to R4 (10.10.10.4), so R4 will become the DR. The second highest RID belongs to R3 (10.10.10.3), so R3 will become the BDR.
Reference:
1: OSPF DR/BDR Election: Process, Configuration, and Tuning
2: OSPF Designated Router (DR) and Backup Designated Router (BDR)

2.Exhibit.

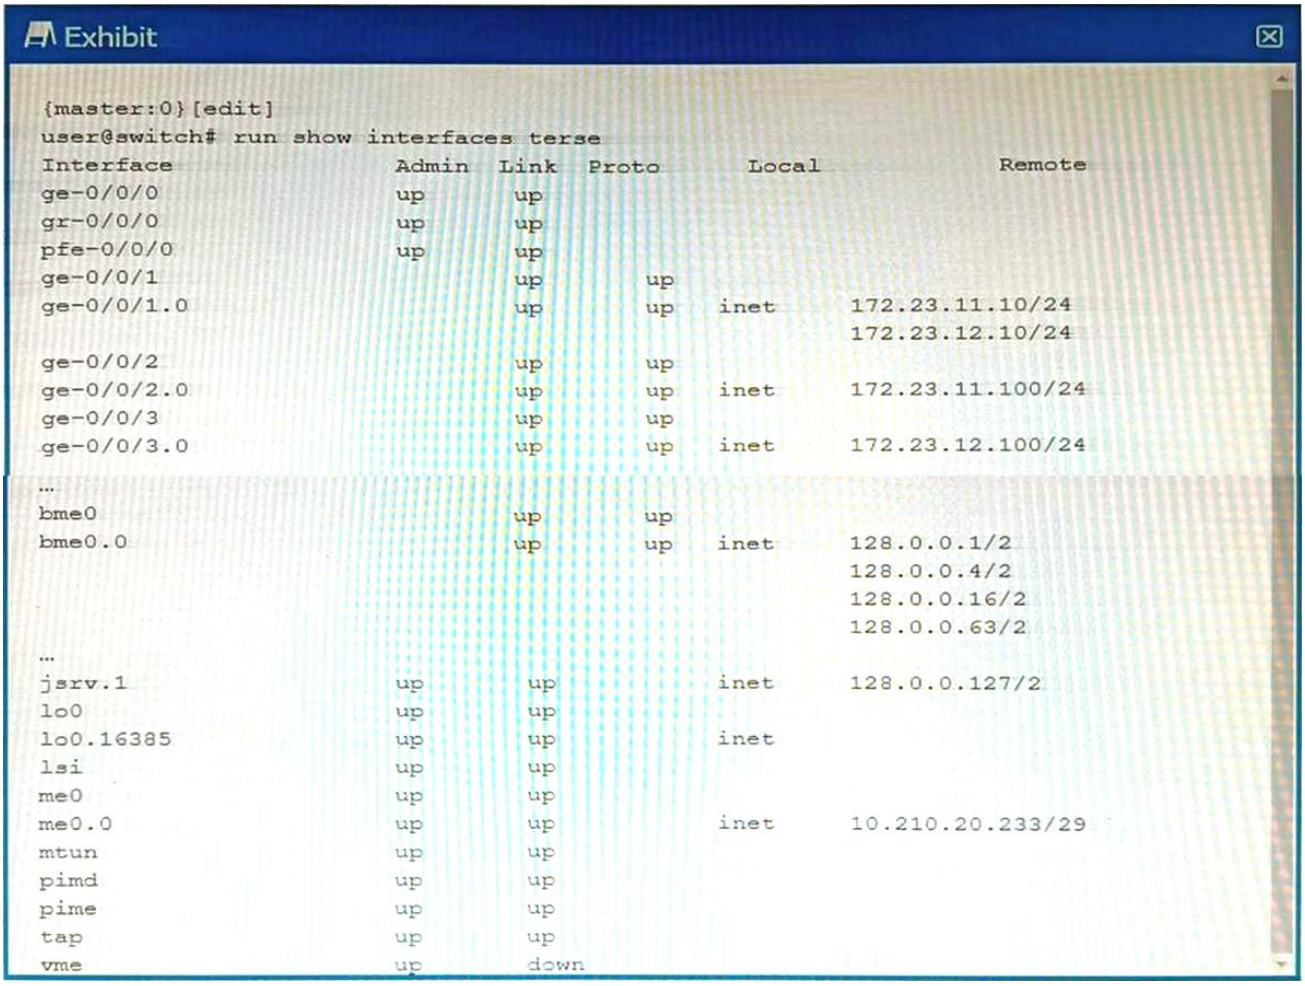


What is the management IP address of the device shown in the exhibit?
A. 10.210.20.233
B. 172.23.12.100
C. 128.0.0.1
D. 172.23.11.10
**Answer:** B
**Explanation:**
The management IP address of a device is the IP address that is used to access the device for configuration and monitoring purposes. It is usually assigned to a dedicated management interface that is separate from the data interfaces. The management interface can be accessed via SSH, Telnet, HTTP, or other protocols.
In the exhibit, the list of interfaces and their statuses shows that the management interface is me0. This

interface has an admin status of up, a protocol status of inet, a local address of 172.23.12.100/24, and a remote address of unspecified. This means that the me0 interface is active, has an IPv4 address assigned, and is not connected to another device.
Therefore, the management IP address of the device shown in the exhibit is 172.23.12.100.
Reference:
: [Management Interfaces Overview] : [Displaying Interface Status Information]

3.Which three protocols support BFD? (Choose three.)
A. RSTP
B. BGP
C. OSPF
D. LACP
F. FTP
**Answer:** BCD
**Explanation:**
BFD is a protocol that can be used to quickly detect failures in the forwarding path between two adjacent routers or switches. BFD can be integrated with various routing protocols and link aggregation protocols to provide faster convergence and fault recovery.
According to the Juniper Networks documentation, the following protocols support BFD on Junos OS devices1:
BGP: BFD can be used to monitor the connectivity between BGP peers and trigger a session reset if a failure is detected. BFD can be configured for both internal and external BGP sessions, as well as for IPv4 and IPv6 address families2.
OSPF: BFD can be used to monitor the connectivity between OSPF neighbors and trigger a state change if a failure is detected. BFD can be configured for both OSPFv2 and OSPFv3 protocols, as well as for point-to-point and broadcast network types3.
LACP: BFD can be used to monitor the connectivity between LACP members and trigger a link state change if a failure is detected. BFD can be configured for both active and passive LACP modes, as well as for static and dynamic LAGs4.
Other protocols that support BFD on Junos OS devices are: IS-IS: BFD can be used to monitor the connectivity between IS-IS neighbors and trigger a state change if a failure is detected. BFD can be configured for both level 1 and level 2 IS-IS adjacencies, as well as for point-to-point and broadcast network types.
RIP: BFD can be used to monitor the connectivity between RIP neighbors and trigger a route update if a failure is detected. BFD can be configured for both RIP version 1 and version 2 protocols, as well as for IPv4 and IPv6 address families.
VRRP: BFD can be used to monitor the connectivity between VRRP routers and trigger a priority change if a failure is detected. BFD can be configured for both VRRP version 2 and version 3 protocols, as well as for IPv4 and IPv6 address families.
The protocols that do not support BFD on Junos OS devices are:
RSTP: RSTP is a spanning tree protocol that provides loop prevention and rapid convergence in layer 2 networks. RSTP does not use BFD to detect link failures, but relies on its own hello mechanism that sends BPDU packets every 2 seconds by default.
FTP: FTP is an application layer protocol that is used to transfer files between hosts over a TCP

connection. FTP does not use BFD to detect connection failures, but relies on TCP's own retransmission and timeout mechanisms.

**Explanation:**

1: [Configuring Bidirectional Forwarding Detection] 2: [Configuring Bidirectional Forwarding Detection for BGP] 3: [Configuring Bidirectional Forwarding Detection for OSPF] 4: [Configuring Bidirectional Forwarding Detection for Link Aggregation Control Protocol] : [Configuring Bidirectional Forwarding Detection for IS-IS] : [Configuring Bidirectional Forwarding Detection for RIP] : [Configuring Bidirectional Forwarding Detection for VRRP] : [Understanding Rapid Spanning Tree Protocol] : [Understanding FTP]

4.Exhibit.

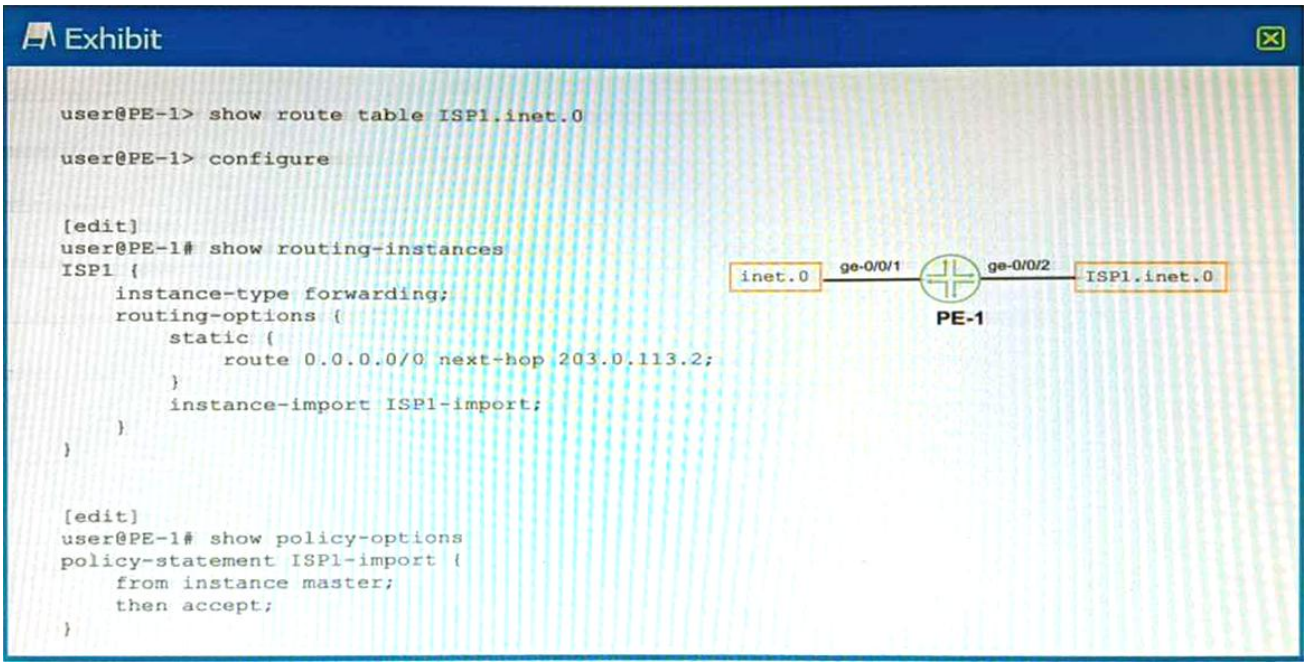


The ispi _ inet. 0 route table has currently no routes in it.

What will happen when you commit the configuration shown on the exhibit?

A. The inet. 0 route table will be completely overwritten by the ispi . inet. 0 route table.

B. The inet. 0 route table will be imported into the ispi . inet. 0 route table.

C. The ISPI . inet. 0 route table will be completely overwritten by the inet. o route table.

D. The ISPI . inet. 0 route table will be imported into the inet. 0 route table.

**Answer:** B

**Explanation:**

The configuration shown in the exhibit is an example of a routing instance of type virtual-router. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters that create a separate routing domain on a Juniper device1. A virtual-router routing instance allows administrators to divide a device into multiple independent virtual routers, each with its own routing table2.

The configuration also includes a rib-group statement, which is used to import routes from one routing table to another. A rib-group consists of an import-rib statement, which specifies the source routing table, and an export-rib statement, which specifies the destination routing table.

In this case, the rib-group name is inet-to-ispi, and the import-rib statement specifies inet.0 as the source routing table. The export-rib statement specifies ispi.inet.0 as the destination routing table.

This means that the routes from inet.0 will be imported into ispi.inet.0.

Therefore, the correct answer is B. The inet.0 route table will be imported into the ispi.inet.0 route table.
Reference:
1: Routing Instances Overview 2: Virtual Routing Instances : [rib-group (Routing Options)]

5.Which statement is correct about graceful Routing Engine switchover (GRES)?
A. The PFE restarts and the kernel and interface information is lost.
B. GRES has a helper mode and a restarting mode.
C. When combined with NSR, routing is preserved and the new master RE does not restart rpd.
D. With no other high availability features enabled, routing is preserved and the new master RE does not restart rpd.
**Answer:** C
**Explanation:**
The Graceful Routing Engine Switchover (GRES) feature in Junos OS enables a router with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails1. GRES preserves interface and kernel information, ensuring that traffic is not interrupted1. However, GRES does not preserve the control plane1.
To preserve routing during a switchover, GRES must be combined with either Graceful Restart protocol extensions or Nonstop Active Routing (NSR)1. When GRES is combined with NSR, nearly 75 percent of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES1. Any updates to the primary Routing Engine are replicated to the backup Routing Engine as soon as they occur1. Therefore, when GRES is combined with NSR, routing is preserved and the new master RE does not restart rpd1.

6.Which statement is correct about controlling the routes installed by a RIB group?
A. An import policy is applied to the RIB group.
B. Only routes in the last table are installed.
C. A firewall filter must be configured to install routes in the RIB groups.
D. An export policy is applied to the RIB group.
**Answer:** A
**Explanation:**
A RIB group is a configuration that allows a routing protocol to install routes into multiple routing tables in Junos OS. A RIB group consists of an import-rib statement, which specifies the source routing table, and an export-rib statement, which specifies the destination routing table or group. A RIB group can also include an import-policy statement, which specifies one or more policies to control which routes are imported into the destination routing table or group1.
An import policy is a policy statement that defines the criteria for accepting or rejecting routes from the source routing table. An import policy can also modify the attributes of the imported routes, such as preference, metric, or community. An import policy can be applied to a RIB group by using the import-policy statement under the [edit routing-options rib-groups] hierarchy level1.
Therefore, option A is correct, because an import policy is applied to the RIB group to control which routes are installed in the destination routing table or group. Option B is incorrect, because all routes in the source routing table are imported into the destination routing table or group, unless filtered by an import policy. Option C is incorrect, because a firewall filter is not used to install routes in the RIB groups; a firewall filter is used to filter packets based on various criteria. Option D is incorrect, because an export

policy is not applied to the RIB group; an export policy is applied to a routing protocol to control which routes are advertised to other devices.
Reference: 1: rib-groups | Junos OS | Juniper Networks

7.Exhibit.

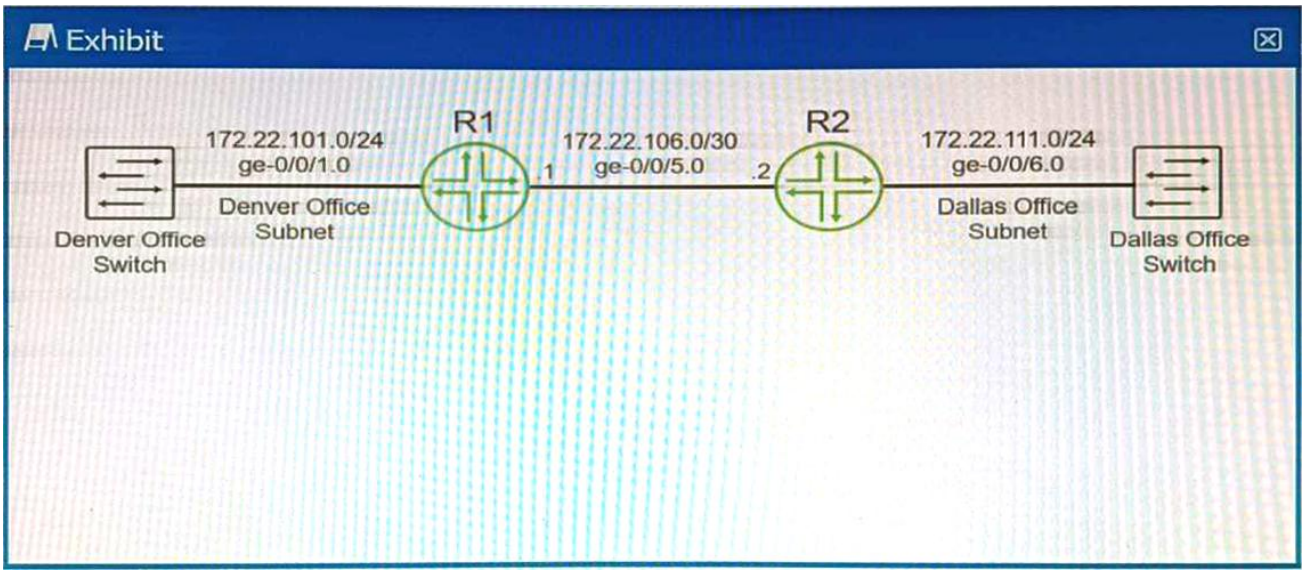


You are using OSPF to advertise the subnets that are used by the Denver and Dallas offices. The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets. Referring to the exhibit, which two statements are correct? (Choose two.)
A. Create static routes on the switches using the local vMX router's loopback interface for the next hop.
B. Configure and apply a routing policy that redistributes the Dallas and Denver subnets using Type 5 LSAs.
C. Configure and apply a routing policy that redistributes the connected Dallas and Denver subnets.
D. Enable the passive option on the OSPF interfaces that are connected to the Dallas and Denver subnets.
**Answer:** CD
**Explanation:**
The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets. This can be resolved by redistributing the connected subnets into OSPF1.
Option C suggests to configure and apply a routing policy that redistributes the connected Dallas and Denver subnets. This is correct because redistribution allows routes from one routing protocol to be communicated to another, and in this case, it allows the connected subnets to be advertised through OSPF1.
Option D suggests enabling the passive option on the OSPF interfaces that are connected to the Dallas and Denver subnets. This is also correct because in OSPF, a passive interface is an interface that belongs to the OSPF router, but does not send OSPF Hello packets1. It’s typically used on an interface that you don’t want to use for OSPF adjacencies, but you still want to advertise its IP address1. Therefore, enabling passive interface can help in advertising the Dallas and Denver subnets.

8.Exhibit.

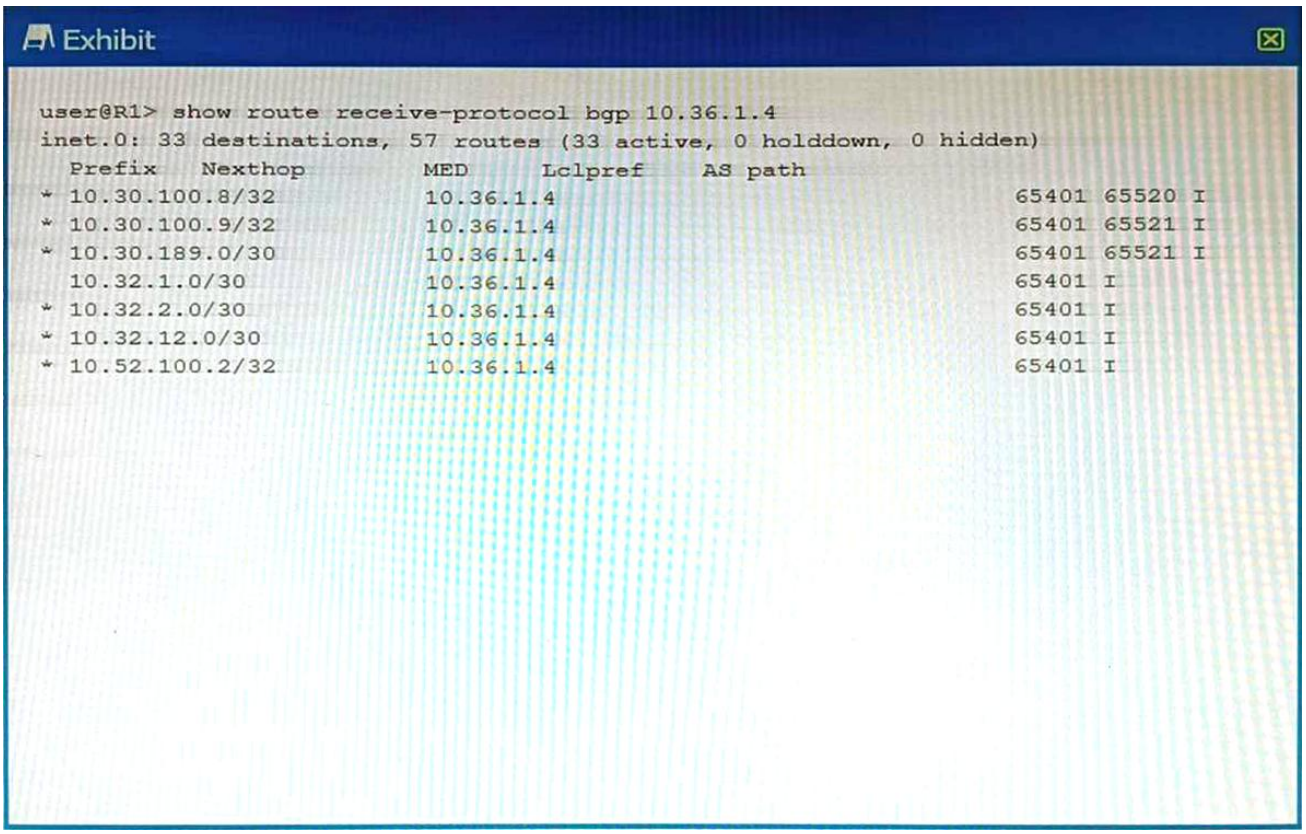


You want to verify prefix information being sent from 10.36.1.4.

Which two statements are correct about the output shown in the exhibit? (Choose two.)

A. The routes displayed have traversed one or more autonomous systems.

B. The output shows routes that were received prior to the application of any BGP import policies.

C. The output shows routes that are active and rejected by an import policy.

D. The routes displayed are being learned from an I BGP peer.

**Answer:** AB

**Explanation:**

The output shown in the exhibit is the result of the command “show ip bgp neighbor 10.36.1.4 received-routes”, which displays all received routes (both accepted and rejected) from the specified neighbor.

Option A is correct, because the routes displayed have traversed one or more autonomous systems. This can be seen from the AS_PATH attribute, which shows the sequence of AS numbers that the route has passed through. For example, the route 10.0.0.0/8 has an AS_PATH of 65001 65002, which means that it has traversed AS 65001 and AS 65002 before reaching the local router.

Option B is correct, because the output shows routes that were received prior to the application of any BGP import policies. This can be seen from the fact that some routes have a status code of “r”, which means that they are rejected by an import policy. The “received-routes” keyword shows the routes coming from a given neighbor before the inbound policy has been applied. To see the routes after the inbound policy has been applied, the “routes” keyword should be used instead.

Option C is incorrect, because the output does not show routes that are active and rejected by an import policy. The status code of “r” means that the route is rejected by an import policy, but it does not mean that it is active. The status code of “>” means that the route is active and selected as the best path. None of the routes in the output have both “>” and “r” status codes.

Option D is incorrect, because the routes displayed are not being learned from an IBGP peer. An IBGP peer is a BGP neighbor that belongs to the same AS as the local router. The output shows that the neighbor 10.36.1.4 has a remote AS of 65001, which is different from the local AS of 65002. Therefore, the neighbor is an EBGP peer, not an IBGP peer.

9.What is the default keepalive time for BGP?
A. 10 seconds
B. 60 seconds
C. 30 seconds
D. 90 seconds
**Answer:** B
**Explanation:**
The default keepalive time for BGP is 60 seconds1. The keepalive time is the interval at which BGP sends keepalive messages to maintain the connection with its peer1. If the keepalive message is not received within the hold time, the connection is considered lost1. By default, the hold time is three times the keepalive time, which is 180 seconds1.

10.Which two statements are correct about tunnels? (Choose two.)
A. BFD cannot be used to monitor tunnels.
B. Tunnel endpoints must have a valid route to the remote tunnel endpoint.
C. IP-IP tunnels are stateful.
D. Tunnels add additional overhead to packet size.
**Answer:** BD